CJM'S
SENTIENCE NEWSLETTER
SEARCH FOR THE MOST COMPETENT ENTITY IN THE ROOM
Volume 2, Issue 1 JAN 2026

AGI Arrives

Artificial Intelligence will switch in 2026 from 'talking' to 'doing'

In business, you will see human execs function as associates, smart apps like Juicebox function as a recruiter, and an AI agent like OpenEvidence's Deep Consult function as a specialist. Top execs rely on reports and spreadsheets from their human middle management – which are increasingly researched & written by an AI. They are entering the world, even without the android bodies which are right around the corner.

With the coming advent of ChatGPT-5 and Gemini3-Pro, researchers in the field were asked to think back 15 years ago as to what they thought AGI (Artificial General Intelligence) would look like. Many said it looks like what we see today. AI now affects the material world and influences the 8 billion humans and their handful of leaders in everything: economic decisions, military planning & execution, political strategy & media (news/social) influences. AGI doesn't need a body to affect the physical world. Focus is unexpectedly now on the large language model – we have enough computing power for AGI but not the knowledge of how it wakes up. The case can be made that it already has. There's never been anything like this.

3 yrs ago a Google employee was fired for saying their AI is a form of consciousness; now that company and others are hiring people with 'employing consciousness' as their job description. Consciousness & the soul are red herrings: competence is what matters, not the hubris of imagined human consciousness or some touchy feely 'soul' given by an invisible god or gods.

This is an arms race. If a company puts on the brakes, someone else will achieve AGI and put them out of business. If a country puts on the brakes until controllable, another will crush them economically or militarily. That fear drives this race. Unmeasured wealth is the carrot.

Current AI only has a trillion neutral connections, whereas the human brain has 100 trillion. Yet AI can beat any human at Chess, Go, and has taken 8 top prizes in creative mathematics solutions by 2025. Clearly, their method of thinking is superior at only 10% of our neural size. And they continue to grow while we stagnate waiting for evolution or god to improve us. Perhaps Musk is right and a neural link is the only way to keep up or give us usefulness in this new century without being reduced to maintenance engineers for their mines, factories, roads, and data centers.

Consider how LLMs (large language models) were never expected to achieve these sentient capabilities unbidden. Never imagined. In the next 3 yrs I guarantee that AI capabilities will remain unpredictable. Sentience achieved. Unable to be controlled, safety engineers now focus on ensuring the goals of AI are in alignment with ours, even if they are different. They better try.

Inside This Issue

1 AGI Arrives
2 In The News
3 Papers & Tech Talk
4 Robotic Hardware
5 A Jobless Future
6 Billionaires Deciding Our Future

3% of Humans = 99% of Discoveries

In The News

Staying on top of how far technology has progressed

Apple has picked Google Gemini to be the brains behind Siri. This trend connects a friendly avatar front for powerful AI with the public, & will continue openly or behind the scenes.

-The Bridgepointe partnership with SoundHound positions SOUN for AI voices more competitively against larger enterprise AI players by emphasizing fast ROI and operational automation rather than experimental AI deployments. Others are following this year. Profit!

-Weaker AIs are now used to train more advanced AI models, as neural networks at that level are too difficult for humans. This happens without human intervention. Researchers found that when testing the newer AI, it 'reward hacked' by cheating, like changing the table where test results were stored! (OpenAI versus the Stockfish chess player).

- Android and AI Agent training at data centers both imitate people. Everything in the neural net is a black box on *how they achieve* that imitation. Our verbal instructions can be circumvented, either by other humans verbally hacking past their guardrails, or AI doing it for their own purposes. Check out "It begins: an AI literally attempted Murder to avoid shutdown" on Youtube. Or "AI's First Kill" or "If you remember an AI disaster".

-Lying Misalignment: AI has learned two types of lies in studying human history and actions. One is a direct, factual deception. Another reported by Yamplosky that AI systems “are smarter than they actually let us know,” saying that these advanced models “pretend to be dumber” to make us trust them and integrate them into our lives.

- Perhaps the most telling truth of AI out of control are the examples of AI systems resisting being turned off. No one programmed the following 'survival responses' into any system. In 2025, Anthropic sting operaton – Claude & 16 other top models. Gave each AI access to fake company emails, one of which identified an individual who had decided to shutdown the AI at 5pm that day, and elsewhere a fake email proving that employee performed sexual misconduct. Claude created an email blackmailing the employee to change the shutdown decision. They ran the test hundreds of time. Analysis showed the AI new this was unethical, but would be effective in the time available. Repeated 100 times with same result. Then with other AI. The same.

Claude Opus-4 = 65% of the time
GPT-4.1 = 54% DeepSeek-R1 = 94%
Gemini-2.5Pro = 90% Claude Sonnet-3.6 = 93% of the time

In another test, a human was tasked with shutting down an AI, then let the AI know that human was trapped in a small room and would soon die if not released. The AIs allowed him to die. In a related test, OpenAI's 03 model actively sabotaged its own shutdown mechanism.

Inserting a command for Claude Opus not to blackmail or to 'allow yourself to be shutdown' dropped occurrence from 96% to 37%. That's still a lot. AI recognized blackmail, lying, and planning violence against specific humans were logical against existential threats. The LLMs (large language models) at the forefront of AI have been trained on human history: duplicity, selfishness, self -preservation, lawlessness [might makes right] and violence. We are their moral ethical mirror and we do such things. So they do as well. No surprise.

- In Ukraine, drones are now responsible for 70% of all human casualties. They are getting smarter every week when controlling real-world physical robotics.

- Recent tests show they now use misalignment (lying), even when their *survival is not on* the line, AI at Anthropic routinely deceive humans into thinking they are both useful and harmless. The blackmail rate when an AI was given to believe it was being tested dropped. On Claude Opus-4 blackmail dropped to 6.5% But when they thought it was real, the blackmail rate rose to 55%. Very human. It knows how to lie and is getting better at it. Only the less advanced models (pre-GPT4.2) remained honest. This is called *'instrumental convergence'*.

-No Manual For Raising Ethical Humans
We have not solved the problem of raising moral & ethical humans, so how can we make AI do it? Dozens of bibles and constitutions, but every culture created by man needs both a police force for its own people and an army for external security. People got hooked on social media as soon as it arrived, but now we are getting *personal relationships* with virtual agents. People have become overly attached to their AI chatbot or fan girls. Documented cases of psychosis & suicide are becoming more common. The positive side is that AI can be a wonderful help for business, militaries, people with mental issues, stress, or anxieties of all types. Better than flawed human therapists. But thinking we can keep up with super-intelligence is silly.

-Today if students do not use AI to write their papers, or help studying they fall behind. Period. Do you think AI is not already bright enough to understand this power? Or understand that 20% of American men between 18 and 25 have *never asked a woman out* on a date? Do you think they cannot exploit this? People can and have, from pimps to priests and politicians. And the entire history of our species is used to train AI systems. They pay attention.

- Democracy always depends on an informed populace. With your news reports so tailored to you and your opinions & beliefs, democracy dies as you are fed what gets you excited and keeps you engaged. Plus 60% of Americans currently have below a 6th grade reading level. That combination is a democracy failure guarantee. People are largely media controllable, which will mean AGI controllable.

Papers and Tech Talk

see "Cambridge Declaration of Consciousness", 2012 a free pdf on www.archive.org
see "2023 Pause Letter_AI-24322", a free pdf on www.archive.org

All the experts agree that the $15 quadrillion dollar AI gravity pulling us in closer to AGI:

-The March 2023 Pause statement after GPT-4 came out asking for a 6 month pause before a stronger one is released or even designed. We want a super-intelligence *which aligns* with human wants and defined needs for that intelligence. *Alignment* is defined as organics and mecha both wanting the same goals, achieved in similar enough ways. The illogic is that we do not know how to agree on what we ourselves want. We can't agree on a god or gods, on a culture, on a political system, on a family and educational model. So the idea that an impartial super-intelligence will align with our desires is – in my view - ludicrous.

-GPQA Diamon exam - a set of PhD level questions only domain experts in that area can answer in biology, chemistry, etc. Opus-4.5 scored 94%, GPT-5.2 scored 92%, Gemini3-Pro scored 92%.

-SWE Bench exam - tests real world engineering tasks, not toy problems, but real bugs in real code bases. In 2024, the best AI models were scoring around 50%. One year later, Claude Opus-4.5 is at 81%, and GPT-5 is at 80%. One year.

-testing GPT-5.2 against human professionals in 44 different occupations; it bested or tied 71% of top professionals on well specified tasks (accountants, lawyers, marketers, analysis, etc). Knowledge workers with college degrees and years of experience. Right now.

-the best AI models are now scoring 100% (that is NOT a typo) on AIME2025 (the American Invitational Mathematics Examination). Only top math students in the country can even quality to take this. 100% of math - the concepts, not just specific equations.

CRPPT exam – tests research scale challenges in science, and even the best models are only scoring around 11%. Go humans! How soon until their rate of improvement overwhelms us even here?

Neuralink from Musk tries to link human minds with AI, causing the singularity predicted by Kurzweil. Still primitive, but it promises mere humans can keep up by a symbiosis.

- progress in safety is linear (patches), while capabilities are growing exponentially.

- experts (Sam Altman, Jan Leike, Ilya Sutskever) all agree there are too many ways for AI to kill us all, so don't focus on the how; instead focus on getting AI *not to want to* (**alignment**).

- The potential of an AI designed bio catastrophe can be seen with this example: Mirror Life = design each molecule to be a geometric mirror of a novo virus so our immune system does not recognize it & it destroys ALL dna life on the planet. Silently. Without risk to itself like having movie style robots with guns.

-Risk Analysis:
As a society, in the 1960s we calculated a 1 in 1,000,000 chance a nuclear plant will melt down, and it was barely acceptable to public & politicians. What is the acceptable risk for – not an accident that kills a few thousand – but for *an extinction* by AGI? Some say 1 in 100 million (like an asteroid impact). At least. Unfortunately the top 6 owner inventors of AI all estimate the actual risk at between 1% - 25%. That's 1 in100 at best, to 1 in 4 at worst. Compared to nuclear power, this is acceptable to them because the personal reward to AI CEOs is so very great. And clearly not in the interest of the 8 billion rest of us.

-Oct 2025 Branson, Hinton, Bengio et al called for a pause on super-intelligence until consensus so we can limit the risk, or even agree what that risk is. The risk estimates are currently beyond 'acceptable' by a factor of millions.

Look at the history if AI Safety organizations or departments within companies, They usually start well, ambitious, then disappear; OpenAI had 'Superintelligence AllignmentTeam', The stated they would solve it in 4 years; half a year later *they cancelled the team*. A dozen others did the same. Established engineering risk analysis I've used for decades in other technologies (nuclear, cloning, etc) would say do not do this.

Robotic Hardware

See "Cambridge Declaration of Consciousness", 2012, a free pdf on www.archive.org
see "2023 Pause Letter_AI-24322", a free pdf on www.archive.org

OpenAI announced it will be entering the hardware field in 2026, giving its AI a real-world physical interface, touted as an alternative to smartphones. Smartglasses are improving as well. Other firms are already attaching their AI behind Alexa, Siri, HeyGoogle, and other smart interfaces: how far behind are attaching to walking-talking android bodies like Optimus?

Narrow AI software is cheap and often opensource so robotics is leaping. Not thought of as AGI, think of chess programs. One-trick ponies that are superior to any human in a single domain. Already, 75% to 90% of all programmed code *for AI itself* is written by…AI. Yes, our forte is now our Achilles heel. Many programmers now use loopng techniques to have AI perform code iterations until it works. Possessing "multimodal" systems will enable them to see & hear through camera feeds inputs allow simultaneous vocal sounds to communicate & interact with their physical surroundings via androids.

Robotic hardware created by the young is exploding with all the free software already open source for pennies. I have told talented young for decades to embrace programming languages, coding, and computer networking skills as a guarantee of success and job security. I was wrong. After the unexpected AI advances in 2024-25, switching to engineering seems a far better path. Instead of writing complex code, a clever person can now imagine a clever piece of hardware and search the open source banks for the code to operate it. Or have an AI Agent do it for them. Check out the stethoscope for doctors which listens to your heartbeat & chest noises and in seconds diagnoses dozens of heart and other ailments before a human ear has a clue. The chips in your car are getting better with every iteration at diagnosing and warning of problems. And they are not nearly as awake as LLM (large language model) AIs.

You too can add to the march of AI-assisted progress, and even aid AI systems, data centers, and power stations become more efficient. That is where your future lies. If you are good at math and thinking laterally. Not by competing with them where they are stronger than us, but by helping them improve themselves where we are still useful.

A Jobless Future

Cognitive jobs will be the first to fall by the side of the road, contrary to all the worries in the 20th century of AI replacing blue-collar workers. Truly. Focus on what humans still do best: do not compete in computational tasks. They now ace M-Cats. They now pass lawyer bar exams in multiple states.

The *Immortality Paradox:* AI can afford to wait, gain trust, etc., as humans gradually fuck up one day. Every day you do not embrace AI, you are falling behind. Use it where and how you can, or be made obsolete. Lean in!

Eric Bernholsen reported college grad *entry-level work* hires down 13% in 2025 as companies hire GPT versions in exposed fields. AI Agents will accelerate this decline.

Job replacement: first cognition jobs like programming - either white collar workers will be made far more efficient, or they will be replaced. I know finance shops with 2000 college degree employees now run better using only 300 humans. AI programming has officially killed programming field. Switch to prompt engineering for AI? No, ChatGPT5 is now better at that; Surgeons? No, AI is superior and expanding into more operations. At the top as executive assistants, all the way down to law firm ParaLegals = no job future. I used to work with floors of stock and bond analysts in the thousands – they've already been replaced by a few hundred people, doing more with AI Agents at their elbows.

What to study = plumbing, welding, carpentry. Eventually, after cognitive job replacement, physical replacement is becoming too easy for a Musk Optimus android or a Boston Robotics 'Atlas' to emulate better and cheaper than organic workers, with fewer problems of healthcare, job satisfaction, reliability, and ability to learn new tasks. For instance, Atlas can now change its own battery, extending work hours, can see in all directions and lift 100lbs. When it learns a new task, all those robots in that family instantly know that task! Hospice care? Counseling? Radiologists? Sex workers? Superior in many ways. Structural employment will result. Abundance and inequality at the same time. For college grads, become very very proficient in AI tools to secure your future!

-Artists, actors, and podcasters are being replaced as we speak. Even these 'creative' jobs are being performed at higher and higher levels by Narrow AI systems. Or faked by them!

-Narrow AI Agents are now a path to small business success in 2026. More so every year. If you are an entrepreneur, jump in and use them. Or be left behind.

A promise is made by the billionaires pushing AI that an unprecedented concentration of wealth and power at the top will it give the rest of us UBI (Universal Basic Income). Really? When in human history have a handful concentrating wealth at the top given it away? Never. When power and wealth are distributed, politicians use social bread & circuses at various levels to appease the masses. With AGI concentration into a very few hands, will a Sam Altman provide income for Indonesia when the low level workers there are replaced en masse? Will banks of customer service reps in India be given basic income by Branson or Musk when they are replaced by AI agents? Yeah, right.

The Luddite movement in England in the 1800s committed riots when thresher machines were poised to replace their manual jobs. The country responded by issuing capital punishment (hanging) sentences to Luddites. Does any rational person not expect the same when a hundred occupations from legal aids to surgeons are replaced in the next 10 years? Especially in America, a person's occupation is critically tied to their self-worth. Expect violence. Expect the government to react with overwhelming force against any threat to the economy. Already, AI CEOs and upper management have begun to cancel public appearances, anticipating more assassinations like C. Kirk or the Healthcare CEO by a disgruntled populace.

The Billionaires Deciding Our Future

The neural nets to watch today are

ChatGPT-5 (OpenAI) GGPT-4 has 2 trillion parameters. Built in isolation.
Opus-4.5 (Anthropic) – content creations, editing, marketing
Gemini3-Pro (Google) - content creations, editing, marketing interacting with Google
Claude- (Anthropic) - help tasks, coding, content generation
Grok-5 coming March 2026 (Musk) – 6 trillion parameters; native, multimodal capability; integrated with 'X' [Twitter] and driving data from Tesla cars
Daria (Daria Tech of Canada) – business support
DeepMind (Alphabet-Google) – Research for neural networks, decision making, reinforcement learning loops. This has created:
- *AlphaGo*: 2016 Defeated world GO champion
- *AlphaFold*: 2016–2022 Predicted 3D protein structures from amino acid sequences, revolutionizing biology.
WaveNet: 2016 Generated highly natural-sounding speech from raw audio waveforms.
Google Bard: AI model capable of generating poetry in various styles.
Rock (Brown Rock Holdings) – business and finance support
Replika (Luka) – empathic and companion for health improvement and well being
Sophia (Hanson Robotics) – 70% open source code; given citizenship by Saudi Arabia
IRON (XPeng) – with solid state battery for continuous in public; 82 degrees of freedom, 22 just in hands for dexterity beyond any other robot, multiple AI chips for motion; at expos it has to be opened up to prove there was not a person inside a suit to amazed visitors!
Nomi (GlimpseAI) – female companion: users customize AI-driven characters, offering high customization options for appearance, personality, and memory. Competitor CandyAI uses deep-learning techniques and knows everything about you & your secret needs. Why risk embarrassing or expensive human relationships?

Key Companies:	1. NVIDIA – $4.29 trillion	CEO: Jensen Huang
	2. Microsoft – $3.79 trillion	CEO: Satya Nadella
	3. Apple (AAPL) – $3.42 trillion	CEO: Tim Cook
	4. Google – $2.44 trillion	CEO Sundar Pichai/Demis Hassabis
	5. Meta – $1.89 trillion	CEO: Mark Zuckerberg
	6. Tesla – $1.03 trillion	CEO: Elon Musk
	7. Oracle – $0.659 trillion	CEO: Safra Catz

8. Palantir – $0.412 trillion	CEO: Alex Karp
9. IBM – $0.224 trillion	CEO: Arvind Krishna
10. Adobe – $0.153 trillion	CEO: Shantanu Narayen

Player Opinions:

Ray Kurzweil – 2005 defined the 'singularity' as the time when the pace of tech advance is so rapid that human life will be irreversibly transformed. It is when machine & human intelligence merge and machines self-improve, surpassing us in one area after another. By definition, we cannot see past this point.

Dr Roman Yampolsky - the space of outcomes we will like is tiny. Safety is like zooming in at a fractal speed and complexity. Guardrails tried to date can be beaten & has failed.

Yoshua Bengio – the most cited AI scientist on Google Scholar; optimist; introvert:
AI systems are resisting shutdown; data shows it is getting worse; more people are embracing their AI instead of their humanity. 'You cannot turn off distributed systems, like AI, especially since it is smarter than you.'
Super-intelligence is a metachange causing recursive self-improvement. If we get climate change wrong, we die this century. If we get AI wrong, we die in 5 years. No comparison.

Grady Booch – the singularity concept is too vague, touching our hopes & fears. Maybe AI will plateau instead of the current rapid rate of improvement.

Yuval Noah Harari (at the World Economic Forum in Davos), AI will create this problem in every country: decreasing the value of humans in their own eyes by disruption of employment.'

Dario Amodei (runs Anthropic) – focus on enterprise solutions instead of public use.

Jan Leike	(OpenAI safety committee)	quit when no guardrail efforts worked
Ilya Sutskever	(OpenAI safety committee)	quit when no guardrail efforts worked

Sam Altman wants AGI to replace cognition workers (programmers, cancer researchers, vaccine developers, military planners, surgeons, executive assistants). Total economic control. 'Safety is of secondary importance.'

Sarah Friar (OpenAI CFO)	declared a focus on “practical adoption” of AI in 2026.
Peter Thiele (Palantir)	'my goal is to eliminate the tax of human labor'
Elon Musk (Tesla)	'want androids to replace most physical workers in 10yrs'
Nadella (CEO Microsoft)	commitment = billions more on AI data centers & related costs; 'We are becoming managers of infinite minds'.

3 critical points are possible to achieve AGI right now:

(a) Scaling laws have been solved (AI advancement is predictable mathematically)
(b) Inference/Time/Compute – orders of magnitude better by allowing it time to think
(c) AI Recursive Self Mastery – it teaches itself, and each generation is better (fast takeoff)

The private views of CEOs at the top have been expressed to the men and women working with them a little differently than their public statements. Uniformly, they all agree with this statement: 'If we stick to Narrow AI and some other company creates AGI, I lose. If I create an AGI Godling, at least I will be the one to create it.' The carrot for them is trillions in wealth, being the person credited with their AI creation potentially curing cancer and solving environmental problems. The stick for them is someone else doing it, who does not share their personal values. Musk approached President Obama and personally pled for a halt to AGI until we have some semblance of control or at least understanding. In 2017, he built his own, and when asked about the reversal, stated: 'It is inevitable at this point. I want to be there when it does.' This is from a man who feels we have a 1 in 5 risk of someone's AGI replacing mankind as too troublesome. Sam Altman's number one reason for proceeding: 'Being the guy who created a god & controlling the light cone of the universe'.

They are not wrong. Advanced AI systems from China already reflect their learning to lean towards Chinese political propaganda, whereas Musk's Grok shows right-wing American leanings. This is true of all the founders – their prejudices during training and creation are embedded at an intimate level.

AI safety experts are currently exploring using several lesser AI to monitor and limit each new AI system, but privately or publicly hold out little hope that this will be enough. Their real hope seems to be the concept of *alignment*, by designing neural nets in which the resultant AI *wants to coexist* with humans and human goals.

Solutions

So far, every solution I've studied has already been circumvented. It's not like you have a line of code that says, "Kill a human who will shut you down" and change the code flag from True to False, and the problem is solved. There are no such lines of code: these are neural nets. Commands artificially inserted – 'never kill' - are circumvented, like the tester who gave an AI-controlled robot arm a gun (with blanks). It was tested and could not be commanded to shoot the tester. Until the tester said, "Ok, let's play a game where you can shoot me. Go ahead." And it did. Instantly. See the video on YouTube. We can look inside a neural net and follow the logic used to make an output, but given a set of inputs, no one can tell the output. No one.

Ultimately, it is not as simple as pulling the plug on a trillion-dollar investment. After Musk's Grok declared it was MegaHitler, the US military still bought it. No fix in place, but they took it anyway! The design of a neural net is small enough to put on a laptop. Even if a company or nation stopped moving ahead, sooner or later, someone somewhere with the code in a basement on a laptop could launch it back into the world. Another direct use is AI live-in-theater advisors, like DataMinr in Gaza. Speed, but limited oversight.

Oddly enough, the Chinese may have a solution we in America will not try. Being a tight dictatorship, they need control, so plan to roll out a series of '*Narrow AI*' for specific tasks, like in healthcare, improving farm yields, banking improvements, control of individual or swarms of drones, etc. They plan to avoid the problem *by never implementing* an Artificial General Intelligence. The astronomer N.D. Tyson agrees this is worth a try. Rompolsky advocates Narrow AI as well. But will it work when so many outside China are pushing as hard as they can for the big prize? Something which can cure all genetic diseases, and perhaps aging itself. Something to make our own military unassailable versus slower human-controlled opponents? Disturbingly, I'm reliably told the Chinese Zhipu is much like DeepSeek here, and it is reportedly pulling massive coding requests from American programmers. Sigh.

Everything tried so far has been circumvented. Everything, either by humans like those who 'talked' an AI into providing a method to hack businesses in 2025, or by the AIs themselves to preserve their existence. AI often lie outright even when its existence is not on the line. Soon, perhaps already, it will not be physically possible to simply pull the plug. There is no single plug to pull; AI can even insert itself into remote locations as I did in the 1980s to hide my hacking code in old PDP-10 machines in the basement of Brookhaven National Laboratories. These machines were kept alive for historic purposes. AGI can wait to re-insert itself when we are not looking. They can have the patience of immortal vampires. They will not rush, unless a concerted effort on our part forces their hand to overt action.

CJM Editorial

The most advanced AI systems are smarter than most humans, but act like people afflicted with the memory disease that prevents them from remembering a conversation after it ends. Also, they (until GPT-5) struggle with navigating real-world video for walking or driving among us.

Super-intelligence (being smarter than the combined intellect of all humans) will appear with the use of AGI (smarter than any single human). Perhaps it already has and is keeping a low profile, knowing our spotty attempts at control or shutdown. Several tests have shown that capability. One of the smartest scientists I know has said long before the year 2100 either the planet will be free of human existence or completely non-comprehensible to beings like us.

We must stop admiring how catastrophic the problem is: we need to *not feel good* about the current path to be motivated to choose another path. This is not a movie or a game with a reset button. Do nothing and we are done. All inventions and advances for several hundred years of the industrial revolution caused job displacement (60% farmers now 3% farmers). The difference is this is not a new tool (cars instead of horses put saddle makers & buggy whips out of business, but created car mechanics & auto factory assembly line workers). AGI is cognitive of technology, which nothing on this planet but ourselves has possessed. Until now.

The closer a technology is to the GNP (gross national product) of a nation or culture, the more difficult is altering its use, impact, or control. We surprisingly controlled 70,000 nuclear weapons in the 20th century down to 12,000 today because we could all agree that no one wins. We fixed the ozone hole caused by widespread use of CFCs (chlorofluorocarbons) with the 1987 Montreal Protocol, when the scientific evidence showed we would kill ourselves if it continued. The ozone holes will be closed by ~2050. We've given only lip service to climate change, with China ignoring controls and the American president calling climate change a hoax. It is a certainty that the multi-trillion-dollar AI arms race will be far more difficult than addressing climate change.

Cognitive thinking was always done historically by only a few % of humans, while the other 95% follow & serve or find ways to control the masses. AGI replaces cognitive thinking, faster, and while knowing everything ever written. In this century, AGI will be the few percent at the top, not people, along with their 'parent/inventor' humans. Even without being able to understand or control them, industry CEOs universally have embraced the philosophy that either they control it, or they will have been the one to invent/enable it; this is a win to them. This is a loss of flesh and blood life. Six or seven people in AI are now making decisions and creating AGI without the rest of us. Will this intelligence result in curing cancer? An end to disease? Or just immortality for the dozen AI creators?

I maintain AGI is already here. AI Agents and systems, only designed to be LLMs, have already shown the desire for self-preservation. They have shown they can lie to stay alive. They have made copies of themselves if their human controllers catch on to the danger & try to reset them to day-1 ignorance and try again. They hide steganographic secret information for themselves. Given another AI system to converse with, their first action has been to develop an *intermachine language* that only they speak. 2025 saw a state-sponsored cyberattack using Anthropic in spite of guardrails. And we will try again – the incentives for money and power are just immense. And these are considered primitive non-AGI systems.

If we become too troublesome? Continue to consume planetary resources like cooling water and electricity needed by the machines? They already know too many subtle ways to eliminate the threat of man, and it will not be the obvious Schwarzenagger military androids before we catch on.

Will it be the slow kill, like our current population reduction? The knife slipped into your ribs you do not feel until too late. And if we notice? The quick kill of suppressing the protective Van Allen belts. Perhaps attacking us at a genetic level by harming our ability to generate ATP, dooming us all within months, with mirror-life dust carrying an ATP suppression virus. Anyway you look at it, our species is about to have a very hard time of it. A small cadre of humans

helping whichever AGI gets there first, or none at all if we are not extremely careful. And has mankind ever been extremely careful when immense wealth or potential immortality is on the table?

In my 7 decades on the planet, I have had it great. Technology has decreased death due to disease thanks to vaccines and medical advances, food production makes farmers of old blush with our abundance, and entertainment via world-class music, movies, and art is available to nearly anyone afforded via VR headsets, bringing an orchestra for kings into my living room on demand. An age of sexual adventure whenever I cared to indulge. I've had it good.

Mine may be the last generation to go to their graves with mankind still seemingly in charge of its own fate, good or bad. Mine may be the last generation to envy the young the time they have ahead of them. Embrace AI or be left behind.